The Recursive Flywheel

Marx, LLMs, and the Game Theory of Cognitive Labor Displacement

---

1. Introduction: What It Means to Work in Technology

The phrase "working in technology" is a euphemism. It sounds neutral, forward-looking, even aspirational—the language of progress, innovation, the future. But it conceals what the work actually is. Stripped of marketing vocabulary, working in technology means one thing: building the apparatus of productivity. It means designing, constructing, deploying, and maintaining the machinery—physical, digital, cognitive—that allows a given quantity of human labor to produce more output than it could before. That is the function. Everything else is branding.

This is not a controversial claim when you trace the history. The weaver's loom, the power loom, the Jacquard loom, the mechanical calculator, the punch-card tabulator, the mainframe, the relational database, the ERP system, the enterprise SaaS platform, the machine learning pipeline, the large language model—every one of these is a tool built by a specialist whose job was to make someone else's job more productive. And "more productive" always has the same meaning in a capitalist economy: the same output with less labor, or more output with the same labor. Which, from the perspective of the laborer being made more productive, means the same thing: fewer of them are needed.

The technology worker is therefore not a neutral participant in the economy. They are a specific character in a specific drama. Marx had a name for the role but no modern term has replaced it: the *producer of constant capital* (Marx, 1976 [1867]). Constant capital, in Marx's vocabulary, is the stock of machinery, tools, and process that sits between the worker and the output. Variable capital is the wages paid to living labor. The distinction matters because the ratio between them—what Marx called the organic composition of capital—is the central variable in the story of industrialization. When constant capital rises relative to variable capital, living labor is displaced. When the opposite happens, labor reabsorbs. In the entire history of capitalism, the opposite has never happened for long.

The technology worker builds the constant capital. That is the job description, translated into the only economic framework that actually makes sense of it. Every line of code, every hardware schematic, every training run, every deployment pipeline, every integration layer is a contribution to the stock of machinery that replaces some quantity of living labor somewhere in the economy. The worker may not see the displaced weaver on the other end of the process. The displaced weaver may work in a different industry,

a different country, a different decade. But the aggregate effect is not in serious dispute. Two and a half centuries of evidence confirm it. The people who build the apparatus of productivity are, in aggregate and over time, the people who displace other people from their livelihoods. This is not a moral judgment. It is a description of what the work is.

What is unusual—what has always been unusual—is that the technology worker has been paid remarkably well for this work. Most labor that builds constant capital has been paid poorly. The machinists who built the power loom were not rich. The foundry workers who cast the steam engine's boilers were not rich. The electricians who wired the first factories were not rich. But the engineers who designed the systems—the small class who translated abstract requirements into working mechanisms—were paid a premium. And for the last fifty years, that premium has been extraordinary. The software engineer, the data scientist, the machine learning researcher, the systems architect, the technical founder—these people have been compensated at rates that put them in the top few percent of the income distribution. Equity compensation has taken some of them further, into the capital-owning class outright. The wages of cognitive labor in technology have been, historically speaking, an anomaly.

The anomaly has a specific explanation. The toolmaking function was scarce. Not "somewhat scarce"—profoundly scarce. It required a combination of abstract reasoning ability, mathematical fluency, years of training, and domain-specific knowledge that only a small fraction of the population could supply. The economic value of a large manufacturing enterprise, a logistics network, a financial system, a telecommunications backbone depended entirely on whether you could find enough people capable of designing and maintaining the machinery. Those people could extract rents because they could not be replaced. The rent-extraction took the form of high salaries, generous stock options, favorable working conditions, and the cultural status of being "in tech." The worker was labor—they sold their labor-power for a wage—but they were labor with the pricing power of capital, because capital's reinvestment strategy (the automation of all other labor) depended on their continued cooperation.

This is worth saying directly because the technology worker's self-understanding has typically obscured it. The dominant self-image of the software engineer, the AI researcher, the technical founder is meritocratic: they earn what they earn because they are smart, because they work hard, because they produce value. The meritocratic story is not wrong, exactly, but it is incomplete. The technology worker's high wages are not simply a return on cognitive ability in some abstract sense. They are a rent on a specific kind of scarcity—the scarcity of people who can translate abstract requirements into the working mechanisms that capital requires to accumulate. Remove the scarcity, and the rent disappears. The cognitive ability remains, but it no longer commands the premium it did.

That scarcity is what created the privilege. And that scarcity is what is now ending. The central argument of this paper is that large language models, specifically, represent the first general-purpose technology capable of performing the toolmaking function itself at scale. The recursive flywheel—the mechanism by which constant capital is produced—no longer requires the same cognitive elite to operate it. The scarcity of the toolmaker is being competed away by the toolmaker's own product. And with the scarcity, the privilege.

This is the historical moment the paper exists to describe. It is a moment that the technology industry has

generally declined to discuss in these terms, for reasons that should be obvious: the people who would have to discuss it are the people whose position is being liquidated, and institutional incentives do not favor clear-eyed discussion of one's own liquidation. But the Marxist framework, supplemented by game theory and several complementary lenses, allows the dynamics to be stated with precision. What follows is that statement.

The argument proceeds in nine further sections. Section 2 presents the Marxist flywheel of surplus extraction and capital reinvestment—the mechanism that has operated continuously for two and a half centuries. Section 3 amends the framework for the specific case of LLMs, which are the first automation technology to target the toolmaker class itself. Section 4 formalizes the dynamics game-theoretically, showing that the Marxist "structural compulsion" is precisely a nested system of multipolar traps whose Nash equilibria systematically diverge from their Pareto optima. Section 5 develops the realization crisis—the core contradiction that capital cannot sell its output to the class whose wages it is destroying. Section 6 turns to the control phase: the game-theoretic necessity of surveillance infrastructure as the mechanism that sustains the equilibrium by eliminating the displaced class's capacity to organize. Section 7 traces the political realignment in which tech capital and traditional capital have formed an alliance against the professional-managerial class whose displacement they are jointly administering. Section 8 presents the two possible long-run equilibria—cybernetic feudalism and the Polanyian counter-movement—and the race condition between them. Section 9 offers supplementary analytical frames (Foucault, Polanyi, Perez, Girard, network theory) that capture dynamics the primary synthesis underweights, Section 10 states what, across all frames, appears to be true, and Section 11 concludes by drawing the argument together and naming the moment the paper exists to describe.

The reader who works in technology should understand that the paper is, in significant part, about them. Not in the trivial sense that everyone is affected by large-scale economic change, but in the specific sense that the technology worker occupies the structural position the paper analyzes: the producer of constant capital whose privileged position depended on a scarcity that is now dissolving. The discomfort of reading this clearly is part of what the reading is for.

2. Theory of the Case: Surplus Value as Engine

Marx's central insight in *Capital* (Marx, 1976 [1867]) is deceptively mechanical. The capitalist does not merely exploit labor—he builds a machine that *accelerates its own capacity to exploit labor*. The process works as follows:

1. **Labor produces surplus value.** The worker produces more value than the wage required to reproduce their labor-power. The difference—surplus value—is captured by the owner of capital.

2. **Surplus value is reinvested in constant capital.** The capitalist takes the extracted surplus and converts it into machinery, tooling, and process—the means of production. This is not optional behavior; it is compelled by competition. The capitalist who does not reinvest is destroyed by the one who does.

3. **Constant capital displaces variable capital.** Each round of reinvestment produces machinery that

replaces some portion of living labor. The loom replaces the weaver. The assembly line replaces the craftsman. The spreadsheet replaces the bookkeeper.

4. **Displacement suppresses the bargaining power of the displaced class.** As machines absorb tasks previously performed by a category of worker, the supply of that labor exceeds demand. Wages compress. The reserve army of the unemployed grows. The remaining workers in that category lose leverage.

5. **Suppressed wages increase the rate of surplus extraction.** Lower labor costs mean higher margins, which means more capital available for reinvestment in the next round of automation.

6. **The cycle repeats.**

This is the flywheel. It is not a one-time event but a self-reinforcing loop: surplus value funds automation, automation displaces labor, displaced labor loses bargaining power, weakened labor produces more extractable surplus, and that surplus funds the next round. Marx called the increasing ratio of constant capital (machines) to variable capital (labor) the *rising organic composition of capital*. He believed it would tendentially depress the rate of profit—but that is a separate argument. The displacement mechanism itself is empirically undeniable and has operated continuously for two and a half centuries.

The critical structural feature: **the flywheel has, until now, always required a privileged cognitive class to design and operate the next generation of constant capital.** The loom displaced the weaver, but someone had to design the loom. The assembly line displaced the craftsman, but someone had to engineer the assembly line. The ERP system displaced the middle manager, but someone had to build the ERP system.

This cognitive class—engineers, programmers, system architects, technical managers—occupied a structurally unusual position in the Marxist framework. They were technically labor (they sold their labor-power for wages), but they were labor whose function was to *build the machinery that displaced other labor*. They were the toolmakers. And because the toolmaking function was scarce—genuinely difficult, requiring rare combinations of abstract reasoning, domain knowledge, and implementation skill—the toolmakers retained extraordinary bargaining power even as they systematically destroyed the bargaining power of every other labor category they touched.

This is the class position of the professional-managerial technical elite from roughly 1970 to 2023. Software engineers, data scientists, systems architects, technical founders. Their wages rose while other wages stagnated. Their equity compensation created a nouveau capital-owning class. Their scarcity was the bottleneck in the flywheel itself—the rate at which capital could automate labor was constrained by the rate at which this cognitive elite could produce the automating systems. Their leverage derived not from solidarity or organization but from the brute fact that the machine could not build itself.

3. The LLM Amendment: When the Flywheel Eats Its Own Operator

Large language models represent a qualitative break in this dynamic. For the first time in the history of the flywheel, the machinery of displacement is capable of performing a significant and rapidly expanding

portion of the *toolmaking function itself*.

This is not a marginal extension of prior automation waves. Prior waves automated *structured, repetitive, well-specified tasks*. They moved from the loom (physical, repetitive) to the robot arm (physical, semi-variable) to the software macro (cognitive, repetitive) to the workflow engine (cognitive, semi-variable). Each wave moved up the ladder of task complexity, but each wave stopped well short of the generative, ambiguous, context-dependent reasoning that characterized the toolmaker's actual work.

LLMs cross that boundary. Not perfectly, not completely, but substantially and with a visible trajectory of improvement. Consider what the toolmaker actually does:

- **Translates ambiguous requirements into structured specifications.** LLMs do this.
- **Produces working code from specifications.** LLMs do this, with increasing reliability.
- **Debugs, refactors, and maintains existing systems.** LLMs do this.
- **Synthesizes information across domains to design novel solutions.** LLMs do this, imperfectly but measurably.
- **Evaluates tradeoffs and makes architectural decisions.** LLMs are beginning to do this.
- **Communicates technical concepts to non-technical stakeholders.** LLMs do this well.

The significance is not that LLMs replace the toolmaker entirely—they do not, yet—but that they *dramatically reduce the scarcity of the toolmaking function*. Where previously you needed a senior engineer with 15 years of systems experience to design and implement an integration pipeline, you now need a moderately technical person with good judgment and an LLM. Where you needed a team of five engineers, you need two. Where you needed rare cognitive ability and a degree from a top-20 program, you need above-average reasoning and the ability to evaluate and direct LLM output.

The effect on the labor market is predictable from the Marxist framework and is already empirically visible:

1. Expansion of the effective labor supply for cognitive work. The number of people who can perform software engineering, data analysis, technical writing, system design, and similar tasks has expanded dramatically. This is the equivalent of what the power loom did to handloom weavers—not elimination, but radical supply expansion that destroys pricing power.

2. Compression of the skill premium. The wage differential between a senior and junior engineer narrows, because the LLM compresses the productivity gap. A junior engineer with Claude or Copilot produces output that, while not equivalent to senior-level work, is much closer than the unaided junior's output was. The senior engineer's premium was based on scarcity; that scarcity is eroding.

3. Acceleration of the flywheel itself. This is the recursive element that makes the LLM wave categorically different. Prior automation waves displaced labor at a rate constrained by the toolmaker class's capacity to build the displacing technology. LLMs accelerate the production of the displacing technology itself. The flywheel no longer has an external speed governor. The rate of automation is now partially a function

of the automation's own output.

4. Shift of value capture from labor to capital. As the cognitive labor that builds AI systems becomes less scarce, the surplus value captured by the owners of the AI infrastructure—the model providers, the compute owners, the platform operators—increases. This is not speculative; it is visible in the market capitalizations of Nvidia, Microsoft, Google, and Anthropic relative to the labor costs of the engineers those companies employ. By early 2026 Anthropic alone had reached a reported private valuation of approximately $380 billion on roughly $14–19 billion of annualized revenue, with capital-to-labor ratios that would have been unrecognizable in any prior technology era (TECHi, 2026).

4. The Game Theory of Structural Compulsion

Marx described the dynamics of capitalist competition with extraordinary precision but without formal tools. Game theory (Nash, 1951; von Neumann and Morgenstern, 1944) provides those tools. What follows is a formalization of the Marxist structural compulsion as a set of nested, interacting games. Readers familiar with Scott Alexander's "Meditations on Moloch" (Alexander, 2014) will recognize the substrate; the analysis here is an attempt to give that intuition a precise game-theoretic and Marxist grounding.

4.1 The Automation Game: N-Player Prisoner's Dilemma

Consider N firms competing in a market. Each firm i chooses a strategy $s_i \in \{\text{Automate}, \text{Don't Automate}\}$. Define:

- $c_i =$ cost of labor for firm i
- $r_i =$ revenue for firm i (function of market share and total market demand D)
- $D =$ total market demand, which is a function of aggregate wages W across all firms and all sectors:

$$D = f(W), \qquad \frac{dD}{dW} > 0$$

This is the critical linkage that Marx identified: demand is endogenous to the wage decisions of firms in aggregate, but exogenous to each individual firm's decision.

The payoff for firm i is:

$$\pi_i = r_i(s_i, s_{-i}, D) - c_i(s_i)$$

The payoff structure is:

	Firm j: Automate	**Firm j: Don't Automate**
Firm i: Automate	$(\pi_{\text{low}}, \pi_{\text{low}})$	$(\pi_{\text{high}}, \pi_{\text{dead}})$
Firm i: Don't Automate	$(\pi_{\text{dead}}, \pi_{\text{high}})$	$(\pi_{\text{med}}, \pi_{\text{med}})$

Where $\pi_{\text{high}} > \pi_{\text{med}} > \pi_{\text{low}} > \pi_{\text{dead}}$.

The Nash Equilibrium is (Automate, Automate). Neither firm can unilaterally deviate to "Don't Automate" without being destroyed. The equilibrium is Pareto-dominated by (Don't Automate, Don't Automate)—both firms would be better off if neither automated, because demand would be preserved. But the cooperative outcome is not a Nash equilibrium; each firm has an incentive to defect.

This is the standard Prisoner's Dilemma structure. Marx's contribution, translated into game-theoretic language: **the Nash equilibrium of capitalist competition systematically diverges from the Pareto optimum, and the mechanism of divergence is labor displacement.**

At N players, the dynamics intensify. Define the fraction of firms that automate as $\alpha \in [0,1]$. Total demand is:

$$D(\alpha) = D_0 \cdot (1 - \lambda\alpha)$$

where λ is the demand destruction coefficient—the rate at which automation erodes purchasing power. Each firm's individual payoff from automating is:

$$\pi_i^{\text{auto}}(\alpha) = \frac{(1+\gamma)}{N} \cdot D(\alpha) - c_{\text{auto}}$$

where $\gamma > 0$ is the competitive advantage from automation (cost savings, market share capture). The individual payoff from not automating is:

$$\pi_i^{\text{no}}(\alpha) = \frac{(1-\gamma\alpha)}{N} \cdot D(\alpha) - c_{\text{labor}}$$

The firm automates if $\pi_i^{\text{auto}} > \pi_i^{\text{no}}$. The critical observation: for any individual firm, automation is the dominant strategy as long as $\gamma > 0$, regardless of α. The competitive advantage of automation is a *local* property; the demand destruction is a *global* property that no individual firm fully internalizes. This is the formal structure of a tragedy of the commons. The commons is aggregate demand. The depleting action is labor displacement. The individual incentive to deplete exceeds the individual cost of depletion at every margin.

4.2 Recursive Acceleration: Iterated Games with Endogenous Tempo

Standard multipolar traps are bad. The LLM variant is worse because the game's tempo is endogenous to its outcomes.

Define the automation capability at time t as $A(t)$, representing the fraction of cognitive tasks that can be automated. In prior technological regimes, $A(t)$ grew at a rate determined by the toolmaker class's capacity:

$$\frac{dA}{dt} = \mu \cdot L_{\text{tool}}(t)$$

where L_{tool} is the labor supply of toolmakers and μ is their productivity. The flywheel's speed was governed by an exogenous constraint—the scarcity of toolmakers.

With LLMs, the automation capability itself contributes to the production of further automation:

$$\frac{dA}{dt} = \mu \cdot L_{\text{tool}}(t) + \beta \cdot A(t)$$

This is a differential equation with a self-reinforcing term. When $\beta > 0$, the growth of $A(t)$ is superlinear:

$$A(t) \sim e^{\beta t} \quad \text{(in the limit as } L_{\text{tool}} \text{ becomes negligible)}$$

The game-theoretic consequence: in each round of the iterated Prisoner's Dilemma, the defector's advantage $\gamma(t)$ is increasing:

$$\gamma(t) = \gamma_0 \cdot A(t)$$

And the cost of cooperation (not automating) is increasing:

$$c_{\text{coop}}(t) = \gamma(t) \cdot \alpha(t) \cdot D(t)$$

This means the Prisoner's Dilemma becomes *more* coercive with each round. The shadow of the future—the factor $\delta \in (0,1)$ representing how much players weight future payoffs—determines whether cooperation can be sustained in a repeated game. The folk theorem (Fudenberg and Maskin, 1986) tells us that cooperation is an equilibrium in infinitely repeated games when δ is sufficiently high:

$$\delta \geq \frac{\pi_{\text{defect}} - \pi_{\text{coop}}}{\pi_{\text{defect}} - \pi_{\text{punish}}}$$

But when π_{defect} is growing exponentially (because automation capability is growing exponentially), the required δ approaches 1—players must care about the future *almost as much as the present* to sustain cooperation. In practice, firms discount the future heavily (quarterly earnings pressure, VC fund lifecycle, competitive urgency), so δ is low and falling. **The recursive acceleration of the flywheel systematically destroys the conditions for cooperative equilibria.**

4.3 Nested Multipolar Traps

The automation Prisoner's Dilemma does not exist in isolation. It is embedded in a system of nested games, each of which reinforces the others.

Trap 1: Firm vs. Firm (the automation race)

As formalized above. Every firm must automate or die. Nash equilibrium is maximum automation.

Trap 2: Nation vs. Nation (the AI arms race)

Consider two nations, US and China, each choosing AI development intensity $d \in [0, d_{\max}]$. The payoff includes both economic competitiveness and military/strategic advantage:

$$U_i(d_i, d_j) = E_i(d_i, d_j) + S_i(d_i, d_j) - K_i(d_i)$$

where E is economic payoff, S is strategic/military advantage, and K is the cost of development. S has a strong first-mover advantage:

$$\frac{\partial S_i}{\partial d_i} \gg 0 \quad \text{when } d_i > d_j$$

The Nash equilibrium is $d_i = d_j = d_{\max}$—maximum development intensity. No nation can unilaterally pause without ceding strategic advantage. This trap *reinforces* Trap 1, because firms that credibly claim national security relevance receive political protection and public funding for their automation activities.

Trap 3: Worker vs. Worker (the skill race)

Individual workers face a symmetric game: adopt AI tools to increase productivity (and compete with peers who have done so) or refuse (and be outperformed). Let $p_i \in \{\text{Adopt}, \text{Refuse}\}$. The payoff structure mirrors the firm-level game:

$$u_i(\text{Adopt}, \text{Adopt}) < u_i(\text{Refuse}, \text{Refuse})$$
$$u_i(\text{Adopt}, \text{Refuse}) > u_i(\text{Refuse}, \text{Refuse}) > u_i(\text{Adopt}, \text{Adopt}) > u_i(\text{Refuse}, \text{Adopt})$$

Each worker who adopts AI tools raises the baseline, depressing the marginal value of all workers' labor. The cooperative equilibrium (collective refusal, preserving the skill premium) is unstable to individual defection. This is the same structure that broke craft guilds: the individual benefit of adopting the power loom exceeded the collective benefit of refusing it, even though universal adoption destroyed the weavers' bargaining power.

Trap 4: Platform vs. Platform (winner-take-all tournament)

AI model providers face an all-pay auction structure. Each firm i invests b_i in compute, talent, and training. The winner (highest capability model) captures market share V. All players pay their bid regardless of outcome:

$$E[\pi_i] = P(b_i > b_{-i}) \cdot V - b_i$$

In an all-pay auction with complete information and N symmetric players, the expected profit for each player is zero in equilibrium—total investment equals total prize value. In practice, this means the AI model providers will collectively spend an amount approaching the total economic value of the AI platform market, with most of that spent by firms that will lose. This is the game-theoretic structure of a speculative bubble, and it is rational for each individual player even though the aggregate outcome is massive value destruction.

Trap 5: Investor vs. Investor (reflexive bubble dynamics)

Capital allocators face a coordination game around AI investment. Define the true value of AI assets as V^* and the market price as P. Each investor i chooses allocation $a_i \in [0, a_{\max}]$. The price P is a function of aggregate allocation:

$$P = g\left(\sum_i a_i\right)$$

Each investor's payoff depends on whether they invested and whether the bubble bursts:

$$u_i = \begin{cases} a_i \cdot (P_{t+1} - P_t) & \text{if bubble continues} \\ -a_i \cdot (P_t - V^*) & \text{if bubble bursts} \end{cases}$$

If the probability of continuation in any given period is q, the investor invests if:

$$q \cdot (P_{t+1} - P_t) > (1-q) \cdot (P_t - V^*)$$

When P is rising rapidly (because other investors are piling in), q can be substantially less than 0.5 and it's still rational to invest. This is the Keynesian beauty contest structure (Keynes, 1936, Ch. 12): each investor is investing not on their assessment of V^* but on their assessment of other investors' assessments of future prices. The game is self-reinforcing and detaches from fundamentals—precisely the dynamic Marx described in his analysis of fictitious capital (Marx, 1981 [1894]).

4.4 The Interaction Structure

These five traps are not independent. They form a reinforcement cycle:

```
Firm automation race (Trap 1)
  → Increases AI capability demand
  → Intensifies platform tournament (Trap 4)
  → Attracts investor capital (Trap 5)
  → Funds more firm-level automation (Trap 1)

National AI race (Trap 2)
  → Provides political cover for firm automation (Trap 1)
  → Funds platform development through defense contracts (Trap 4)
  → Protects investor sentiment via national security framing (Trap
5)

Worker skill race (Trap 3)
  → Increases effective labor supply for AI development
  → Accelerates firm-level automation capability (Trap 1)
  → Feeds data and usage into platform improvement (Trap 4)
```

The system of traps is a fixed point: the equilibrium of each game reinforces the equilibrium of the others. No single trap can be escaped without simultaneously escaping all of them, because deviation in any one game is punished not only within that game but by the dynamics of the adjacent games.

5. The Realization Crisis: The Flywheel's Leak

The flywheel has a structural flaw that Marx identified and that the game-theoretic formalization illuminates. Capital *extracts* surplus value at the point of production, but it *realizes* surplus value at the point of sale. And the buyers are, overwhelmingly, the same wage laborers whose bargaining power the flywheel is systematically destroying.

This is the core contradiction: the Nash equilibrium of the automation game suppresses wages, but aggregate demand is a function of aggregate wages. The firms are depleting the commons (purchasing power) while individually behaving rationally.

Formally, define the aggregate surplus S extracted from labor displacement:

$$S(t) = \int_0^t [Y(\tau) - W(\tau)]\, d\tau$$

where Y is total output and W is total wages. The flywheel increases S by increasing Y (productivity) and decreasing W (displacement). But total revenue R available to realize that surplus depends on demand:

$$R(t) = D(W(t), C(t), G(t))$$

where D is demand as a function of wages W, credit C, and government transfers G. The contradiction is that S and R move in opposite directions as the flywheel accelerates:

$$\frac{dS}{dt} > 0, \qquad \frac{dR}{dW} > 0, \qquad \frac{dW}{dt} < 0$$

Surplus is rising. Wages are falling. Demand, absent external support, is falling. The system is generating more value it cannot sell.

5.1 Historical Patches: Temporary Equilibria

Capitalism has historically managed this contradiction through mechanisms that sustain D even as W declines:

Credit expansion increases $C(t)$ to offset declining $W(t)$. This is a temporal displacement of the contradiction—it does not resolve it, it defers it with interest. The intertemporal budget constraint still binds:

$$\sum_{t=0}^{T} C(t) \leq \sum_{t=0}^{T} W(t) + \text{collateral}$$

When cumulative credit exceeds the discounted value of future wages, the system corrects. This is the formal structure of the 2008 crisis.

Government redistribution sets $G(t) = \tau \cdot S(t)$, taxing some fraction of surplus and transferring it to labor as demand-sustaining income. This works within limits but is a Stackelberg game between capital and the state: capital moves first (choosing automation levels), the state responds (choosing tax and transfer levels), and capital responds to the state's response (relocation, lobbying, regulatory capture).

The subgame perfect equilibrium typically involves less redistribution than the social optimum, because capital's first-mover advantage in the Stackelberg structure allows it to credibly threaten exit.

Financialization replaces the commodity realization of surplus value with the financial realization—asset appreciation, derivatives, rent extraction. This detaches S from R entirely: surplus is "realized" not by selling commodities to wage earners but by inflating asset values that other capital-holders purchase from each other. This is Marx's concept of fictitious capital, and it can sustain the system indefinitely *as long as the fiction is maintained*. The fiction collapses periodically (1929, 2000, 2008), but is always rebuilt because the alternative—confronting the realization crisis directly—is politically intolerable.

5.2 Why the LLM Wave Intensifies the Realization Crisis

Prior automation waves displaced labor categories that consumed primarily necessities—food, housing, basic manufactured goods. The consumption floor compressed but did not collapse, because transfer payments and credit kept it propped up. The system could absorb the displacement of factory workers because factory workers were not the high-margin consumers.

The LLM wave displaces the *professional-managerial class*—the high-consumption, high-margin class. These are the people who sustain the premium end of the demand curve: professional services, premium real estate, SaaS subscriptions, private education, high-end consumer goods. In a post-industrial economy, surplus value realization is disproportionately concentrated in this class's consumption.

The total addressable market for this displacement is staggering. Estimates of global knowledge-worker compensation range from roughly $23 trillion annually for narrower definitions of white-collar labor up to $50–70 trillion when broader definitions of information work are used—approximately 50–70% of global GDP (Forrester, 2018; Manyika et al., 2013). The Institute for Public Policy Research has estimated that roughly 70% of tasks performed in white-collar roles could be "transformed" or "replaced" by generative AI (IPPR, 2024). Even a conservative 30% productivity displacement across this base—replacing three workers with one worker and an AI system—represents a transfer on the order of $7–20 trillion annually from labor income to capital income, depending on which definition of the base is used. This is a surplus extraction event with no precedent in scale, occurring within the class that has historically sustained premium consumption.

The demand consequences are not speculative; they are arithmetic. Define the professional-class consumption multiplier as m_p—the ratio of total economic activity supported by each dollar of professional-class income (including their spending's second- and third-order effects). Empirical estimates of m_p range from 1.5 to 2.5. A multi-trillion-dollar reduction in professional-class income, phased over a decade, implies a proportionally larger reduction in total demand once the multiplier is applied. This is a contraction potentially larger than the 2008 financial crisis by an order of magnitude, phased over a decade rather than occurring as a shock—and unlike 2008, it does not self-correct, because the suppressed wages are the permanent equilibrium of the new automation regime rather than a temporary dislocation of an otherwise functioning labor market.

6. The Control Phase: Surveillance as the Next Capital Cycle

Every prior wave of mass displacement has produced a corresponding investment in the *management* of the displaced. This is not conspiracy; it is a structural response to a game-theoretic problem. When the flywheel generates a surplus population—a reserve army of labor too large to be absorbed by new industries at comparable wages—the system faces a coordination problem: the displaced population has both the incentive and (in the case of the professional class) the capability to organize for redistribution. This organization, if successful, would reduce the rate of surplus extraction—which, from capital's perspective, is a defection from the equilibrium that serves its interests.

Capital therefore invests in *mechanism design*: the construction of institutional arrangements that make acquiescence the displaced class's dominant strategy.

6.1 Historical Precedent

The English Poor Laws and the workhouse system managed the displacement caused by enclosure and early industrialization. The Pinkerton Detective Agency and the National Guard managed the labor unrest of the Gilded Age. The FBI's COINTELPRO managed the political mobilization of the displaced and marginalized in the 1960s and 1970s (Churchill and Vander Wall, 1988). The American criminal justice system—which the Brennan Center estimates consumes approximately $270 billion in annual direct expenditure (Brennan Center, 2024), with the Prison Policy Initiative placing the fuller cost to the government and to families of justice-involved people at approximately $182 billion (Wagner and Rabuy, 2017), and broader social-cost accountings reaching upward of $1 trillion (McLaughlin et al., 2016)—manages the surplus population created by deindustrialization and the collapse of working-class wages since 1975.

In each case, the pattern is the same: displacement generates instability, instability threatens the conditions for continued surplus extraction, and capital responds by funding the infrastructure of social control—which itself becomes a profitable industry and a site of surplus extraction.

6.2 The Mechanism Design of Digital Social Control

The current generation of social control infrastructure is digital. Its formal structure is a mechanism design problem: construct a game Γ' in which the displaced population's dominant strategy is compliance, where the original game Γ (without surveillance) had an equilibrium in which the displaced organize.

In the original game Γ, a displaced worker i chooses between Organize (O) and Comply (C). The payoff depends on the fraction of the population that organizes, α_O:

$$u_i(O, \alpha_O) = \begin{cases} B \cdot P(\alpha_O) - k & \text{if organizing succeeds (probability } P(\alpha_O)) \\ -k - p & \text{if organizing fails} \end{cases}$$

$$u_i(C, \alpha_O) = w_{\text{sub}}$$

where B is the benefit of successful redistribution, $P(\alpha_O)$ is the probability of success (increasing in α_O), k is the personal cost of organizing, p is the punishment for failed organization, and w_{sub} is the subsistence

wage under compliance.

Without surveillance, the organizing game has a tipping-point structure: if α_O exceeds some threshold α^*, organizing succeeds with high probability and the rational strategy is to organize. If $\alpha_O < \alpha^*$, it fails and the rational strategy is to comply. This is a coordination game with two stable equilibria—mass organization and mass compliance—separated by a threshold.

The surveillance apparatus transforms this game by:

1. Increasing k (cost of organizing). Surveillance makes organizing visible, which allows preemptive intervention—firing, financial restriction, legal harassment. The expected cost of organizing rises:

$$k' = k + \sigma \cdot c_{\text{detect}}$$

where σ is the surveillance coverage (probability of detection) and c_{detect} is the cost imposed on detected organizers.

2. Increasing p (punishment for failure). When organizing fails and participants are identified, the punishment can be targeted and severe. Surveillance makes identification near-certain:

$$p' = p + \sigma \cdot c_{\text{punish}}$$

3. Decreasing $P(\alpha_O)$ (probability of success). Surveillance enables targeted disruption of organizing efforts before they reach critical mass—identifying leaders, mapping networks, intervening at coordination bottlenecks:

$$P'(\alpha_O) = P(\alpha_O) \cdot (1 - \sigma \cdot \varepsilon)$$

where ε is the disruption effectiveness.

4. Raising the tipping point α^*. The combined effect is to raise the threshold fraction needed for organizing to succeed. If the original threshold was $\alpha^* = 0.1$ (10% of the displaced population needed to organize), the surveilled threshold might be $\alpha^{*\prime} = 0.4$ or higher—making the coordination problem much harder to solve.

When surveillance is comprehensive ($\sigma \to 1$), the organizing equilibrium can be eliminated entirely, leaving compliance as the unique Nash equilibrium. This is the formal content of the "Sauron" thesis: a surveillance system with sufficient coverage transforms a two-equilibrium coordination game into a single-equilibrium dominance game in which the dominant strategy is acquiescence.

6.3 Palantir as Paradigm

Palantir Technologies is the paradigmatic firm in this space. Founded in 2003 by Peter Thiel, Alex Karp, Stephen Cohen, Joe Lonsdale, and Nathan Gettings, and given its early funding by In-Q-Tel (the venture-capital arm of the Central Intelligence Agency), Palantir's core products are data-integration and analytics platforms that synthesize information across disparate systems—financial records, communications metadata, geolocation data, social media, biometric databases—into a unified analytical environment (Palantir Technologies, 2026; Quartr, 2025). Its two principal platforms are Gotham (developed for defense,

intelligence, and law-enforcement use) and Foundry (developed for commercial and civil-government use), supplemented more recently by Apollo (software delivery) and AIP (an LLM-integrated analytics layer) (Built In, 2026; Palantir Technologies, 2026). Its clients include the CIA, NSA, FBI, DHS, Department of Defense, ICE, all branches of the US military, the UK Ministry of Defence, the Ukrainian Armed Forces, and—as of 2025—the Trump administration's restructured immigration enforcement apparatus through the $30 million *ImmigrationOS* contract (American Immigration Council, 2025). Since the beginning of the second Trump administration in January 2025, Palantir has received over $900 million in new federal contracts (American Immigration Council, 2025).

Peter Thiel, Palantir's founder, named the company after the seeing-stones of J. R. R. Tolkien's *The Lord of the Rings*—the *palantíri*, stones that allowed their holders to see and communicate across vast distances, and which in Tolkien's mythology were corrupted by Sauron for surveillance and domination (Quartr, 2025). The literary reference is precise: this is infrastructure for omniscient surveillance. The capability it provides is exactly what the mechanism design analysis specifies—the ability to detect organizing, identify participants, map networks, and enable targeted intervention. It shifts σ toward 1.

The scale of the white-collar labor pool now subject to automation is relevant here in two directions:

First, it represents the scale of the displacement: tens of millions of educated, networked, politically engaged professionals losing their economic position over a 10–15 year window. This is a population that knows how to organize, how to litigate, how to lobby, and how to run for office. The historical precedent for what happens when a large, educated class is rapidly disenfranchised is not encouraging—the Weimar Republic's professional class, the Soviet Union's displaced intelligentsia, the Arab Spring's underemployed graduates.

Second, the displaced labor compensation represents the scale of the *prize*. The surplus value currently flowing to white-collar labor as wages is, from capital's perspective, value left on the table. Capturing even a fraction of it through automation justifies an enormous investment in the social management of the consequences. If the flywheel can extract several trillion dollars a year in displaced professional wages, it is economically rational to spend a substantial fraction of that on the infrastructure that prevents the displaced from taking it back through political action. The return on investment in social control infrastructure is extraordinary, which is why the market will provide it.

6.4 The Architecture of Digital Social Control

The components of the surveillance mechanism are already deployed and being integrated:

Predictive policing and social scoring. Systems that identify individuals likely to engage in disruptive behavior before they do so. Currently deployed in law enforcement; readily extensible to political organizing, labor action, and regulatory advocacy. These systems increase σ (probability of detection) and decrease the time between organizing activity and state response.

Financial surveillance and control. The ability to monitor, flag, and restrict financial transactions in real time. The infrastructure exists in the banking system's compliance architecture (KYC/AML) and is being extended through CBDC research and cryptocurrency regulation. A population that cannot move money cannot organize. This directly increases k (cost of organizing) by restricting the financial infrastructure of

collective action.

Information environment management. Algorithmic control of what populations see, read, and discuss. Social media platforms already function as attention-allocation systems; the transition from commercial optimization (maximize engagement) to political optimization (minimize destabilizing discourse) requires only a change in the objective function, not a change in the architecture. This decreases $P(\alpha_O)$ by disrupting the information channels through which coordination occurs.

Automated administrative state. AI-driven systems that adjudicate benefits, evaluate applications, score risks, and make decisions at scale without human review. These systems can manage a surplus population with minimal state employment—fewer social workers, fewer bureaucrats, fewer points of human discretion where empathy or solidarity might intervene. This lowers C_{control}, making the control apparatus cheaper and therefore more attractive relative to redistribution.

The synthesis of these components produces a system of social management that is qualitatively different from prior control regimes. It is not violent in the way the Pinkertons were violent. It does not require mass incarceration (though that infrastructure remains available). It operates through friction, visibility, and preemption: making it harder to organize, easier to monitor, and possible to intervene before collective action reaches critical mass.

7. The Political Realignment: Tech Oligarchs, the GOP, and the War on the PMC

The political economy of this transition is not occurring in a vacuum. It is producing—has already produced—a visible realignment of class interests in American politics, analyzable as a coalition-formation game.

7.1 The Old Alignment (1990–2020)

For three decades, the professional-managerial class (PMC)—a category whose classic formulation is due to Ehrenreich and Ehrenreich (Ehrenreich and Ehrenreich, 1977)—was the core constituency of the Democratic Party. College-educated professionals—lawyers, consultants, academics, tech workers, government administrators, journalists, healthcare professionals—provided the Democrats with both votes and cultural hegemony. The party's platform reflected their interests: meritocratic credentialism, regulatory expansion (which employed them), globalization (which benefited them), and social liberalism (which reflected their values).

The Republican Party represented an alliance of traditional capital (finance, energy, manufacturing, real estate) and the white working class, held together by cultural conservatism and the promise that deregulation and tax cuts would produce broadly shared prosperity.

The tech industry sat ambiguously in this alignment. Its workforce was culturally PMC and leaned Democratic. Its founders and capital-holders had class interests that aligned with neither party cleanly—they wanted deregulation (Republican), immigration (Democratic), low taxes (Republican), social liberalism (Democratic), and above all, to be left alone.

7.2 The New Alignment (2020–Present)

The LLM wave has catalyzed a realignment. Model it as a coalition-formation game with four players: Tech Capital (T), Traditional Capital (K), the PMC (P), and the Working Class (W). Each player has a policy preference vector and seeks to form a minimum winning coalition.

Tech Capital's preference: Maximize automation speed, minimize regulation, minimize redistribution, maintain access to global talent, build surveillance infrastructure (which is also its product).

Traditional Capital's preference: Low taxes, deregulation, cheap labor (overlaps with T on most dimensions, diverges on immigration and the specific form of deregulation).

PMC preference: Regulate AI (preserve employment), maintain credentialist gatekeeping, fund the administrative state (which employs them), redistribute (which sustains their clients and their own safety net).

Working Class preference: Job security, wage protection, restrict immigration (reduces labor competition), ambivalent on AI (it hasn't displaced them *yet*, and the PMC's credentialist gatekeeping has been a class enemy).

The old coalition structure was $\{T \cup P\}$ (Democrats) vs. $\{K \cup W\}$ (Republicans), with T as a swing player. The new structure is $\{T \cup K\}$ vs. $\{P\}$ with $\{W\}$ as a contested swing.

The realignment occurs because T's class interest has shifted: the LLM wave makes T's primary project the displacement of P, which means T and P are now in a zero-sum game. T needs K's political infrastructure (the Republican Party) to prevent P from using the state to regulate the displacement. K is happy to provide it, because AI deregulation also serves K's interest in reducing labor costs across their own industries.

The alliance is cemented by several converging interests:

Deregulation of AI. The PMC-Democratic coalition's instinct is to regulate AI deployment—safety requirements, impact assessments, worker protections, algorithmic accountability. Tech capital needs to deploy fast, before the regulatory window closes. The Republican Party's anti-regulatory ideology serves this need.

Reduction of the administrative state. The federal bureaucracy is the PMC's institutional base—it employs them, it is staffed by them, it exercises the credentialist gatekeeping that sustains their class power. DOGE—the Department of Government Efficiency, established by executive order in January 2025 and initially led by Elon Musk—is not merely cost-cutting (Executive Order 14158, 2025). It is the automated displacement of government professionals, using AI tools, directed by tech capital, with the explicit political goal of destroying the institutional base of the opposing class. In game-theoretic terms, it is a move to eliminate P's capacity to play the regulation game by eliminating P's institutional position within the state.

Tax and fiscal policy. Capturing several trillion dollars annually in displaced white-collar labor value and redirecting it to capital requires a tax regime that does not claw it back through redistribution. The Republican commitment to low capital-gains taxes, corporate tax cuts, and the erosion of the social safety

net serves this purpose directly.

7.3 The Chicken Game: PMC vs. Tech Oligarchy

The confrontation between the tech oligarch coalition and the displaced PMC has the structure of a game of Chicken (Hawk-Dove):

	PMC: Resist	**PMC: Capitulate**
Tech: Press	$(-C_T, -C_P)$	$(V_T, -L_P)$
Tech: Accommodate	$(-L_T, V_P)$	(S_T, S_P)

where $V > S > 0 > -L > -C$ for both players.

In Chicken, neither player wants to be the one who swerves, but both prefer swerving to mutual destruction. The equilibrium depends on credible commitment—the player who can convincingly commit to *not swerving* forces the other to capitulate.

DOGE and the broader tech-oligarch political project are commitment devices: by firing federal workers, defunding agencies, and restructuring the state before the PMC can respond, the tech coalition is attempting to make its "Press" strategy irreversible—to convert the sequential game into one where the PMC's choice set is reduced to $\{\text{Capitulate}\}$.

The PMC's counter-strategy would be institutional entrenchment: legal challenges, regulatory moats, bureaucratic resistance, and alliance-building with the Working Class (offering labor protections and redistribution in exchange for political support). Whether the PMC can execute this strategy fast enough—before its institutional base is dismantled—is the central political question of the transition.

7.4 The PMC's Structural Weakness

The PMC's fundamental problem is that its power was always derivative. It did not own capital. It did not control the means of production. It controlled *access*—to credentials, to regulatory processes, to institutional decision-making. This is a form of monopoly rent, but it is a fragile one: it depends on the continued existence of the institutional structures that require credentialed gatekeepers.

When the toolmaker's function is automated, the credentialing and gate-keeping functions that sustained PMC power lose their rationale. You do not need a Harvard MBA to direct an AI agent. You do not need a JD to have an LLM draft a contract. The meritocratic-credentialist order that justified PMC privilege is being undermined by the technology itself.

In game-theoretic terms, the PMC's bargaining position depends on its threat point—what happens if the PMC refuses to cooperate with the new order. If the PMC's skills are scarce and non-automatable, its threat point is strong (the system needs them). If its skills are automatable, its threat point is weak (the system can replace them). The LLM wave is systematically weakening the PMC's threat point, which is why the Chicken game is trending toward PMC capitulation.

8. The Endgame: Cybernetic Feudalism or Polanyian Counter-Movement

The analysis produces two possible long-run equilibria.

8.1 Cybernetic Feudalism: The Stable Dystopia

If the control apparatus deploys successfully—if $\sigma \to 1$, the organizing equilibrium is eliminated, and the displaced class's dominant strategy is compliance—the system converges to a stable low-demand equilibrium. A small capital-owning class extracts rent from an automated economy. A large surplus population is maintained at subsistence through algorithmic management and targeted transfer payments (enough to prevent starvation and riot, not enough to fund organizing). Demand is sustained not by mass wages but by capital-to-capital transactions, state spending funded by taxes on automated production, and a financialized asset economy.

This is not a crisis state. It is an equilibrium—stable, self-reinforcing, and resistant to perturbation. It has historical precedent in feudalism itself: a small landowning class, a large subsistence peasantry, a thin clerical class managing the institutional infrastructure, and enough surplus extraction to sustain the system indefinitely. The digital variant substitutes algorithms for clerics and data for land, but the structure is isomorphic.

The game-theoretic condition for this equilibrium: the surveillance apparatus must be comprehensive enough to eliminate the organizing equilibrium before the displaced class reaches the tipping point α^*. This is a race between the deployment of the control infrastructure and the political mobilization of the displaced.

8.2 Polanyian Counter-Movement: The Democratic Response

Karl Polanyi argued in *The Great Transformation* (Polanyi, 1944) that when markets disembed too far from social relations—when land, labor, and money are treated as pure commodities—society generates a counter-movement to re-embed the economy in social constraints. This is not a Marxist class revolution; it is a broader, messier, cross-class political reaction.

Fascism, the New Deal, and the European welfare state were all Polanyian counter-movements—different in character but all responses to the same disembedding crisis. The LLM displacement could trigger a counter-movement that takes various forms: UBI funded by automation taxes, public ownership of AI infrastructure, antitrust action against model providers, mandatory human-in-the-loop requirements for critical decisions, or outright prohibition of certain AI deployments.

The game-theoretic condition for this equilibrium: the displaced class (PMC + elements of the working class) must form a coalition and achieve political power before the control infrastructure eliminates the organizing equilibrium. This requires:

$$\alpha_O(t^*) > \alpha^{*\prime}(t^*) \quad \text{at some time } t^* \text{ before } \sigma(t^*) \to 1$$

That is: the organizing fraction must exceed the surveillance-adjusted tipping point at some time before the surveillance apparatus is complete.

8.3 The Temporal Structure

Both equilibria are plausible. Which one obtains depends on a race condition:

Rate of control deployment: How fast is $\sigma(t)$ approaching 1? This is a function of AI capability growth, government procurement speed, political will, and legal constraints. Current trajectory: fast and accelerating. DOGE, Palantir's expanding government contracts, and the integration of AI into law enforcement and financial surveillance all increase σ.

Rate of political mobilization: How fast is $\alpha_O(t)$ approaching α^*? This depends on the PMC's ability to organize, form cross-class coalitions, and achieve political power. Current trajectory: slow and fragmented. The Democratic Party has not articulated a coherent AI displacement platform. The PMC is culturally allergic to labor-movement-style collective action. The Working Class does not yet see the PMC's problem as its own.

The buildout phase (2025–2032) is the critical window. During this phase, the displacement is occurring but is masked by investment-driven growth. The control infrastructure is being built but is not yet comprehensive. The displaced class is beginning to feel the effects but has not yet reached the point of political mobilization. Both equilibria remain accessible. After this phase, the path-dependent dynamics of the surveillance-adjusted game are likely to lock in one outcome or the other.

9. Supplementary Analytical Frames

The Marxist game-theoretic synthesis presented above is the primary analytical structure. However, several complementary frameworks capture dynamics that the Marx + game theory lens underweights or misses.

9.1 Foucault: Governmentality and Biopower

The surveillance section of this analysis is more Foucauldian than Marxist. Marx's model of social control is instrumental—the state as "committee for managing the common affairs of the bourgeoisie," wielding repression against recalcitrant labor. Foucault's concept of *governmentality* (Foucault, 2007) is more precise about how modern control actually operates: not through force but through the internalization of discipline (Foucault, 1977), the management of populations through statistical knowledge, the production of subjects who govern themselves.

The algorithmic management system described above—predictive analytics, financial surveillance, information environment control, automated administrative state—is biopower, not class repression. It does not need to imprison the displaced PMC. It needs to make them *legible*, predictable, and self-managing. The mechanism is not punishment but *nudge*: shaping the choice architecture so that compliance is not merely the dominant strategy but the only strategy that feels natural.

Foucault adds a dimension that the game-theoretic formalism misses: the displaced class may not experience its own acquiescence as the result of a strategic calculation. It may experience it as *normal*—as the natural order of things. This is the deepest form of control, and it has no representation in the payoff matrices above.

9.2 Polanyi: The Double Movement

Polanyi's *Great Transformation* (Polanyi, 1944) provides the critical corrective to the Marxist fatalism that the flywheel implies. Polanyi argued that the counter-movement is not optional—it is structural. When market disembedding reaches a critical threshold, society reacts, and the form of the reaction is the central political question of the era.

Polanyi would note that every prior technological revolution produced a counter-movement that eventually re-embedded the economy: the Factory Acts, the New Deal, the welfare state. The question is not whether a counter-movement will occur but *what form it will take*—democratic (New Deal) or authoritarian (fascism). Both are Polanyian counter-movements; both re-embed the economy in social constraints; they differ in who controls the re-embedding and in whose interest.

9.3 Carlota Perez: Technological Revolutions and Financial Capital

Perez's framework (Perez, 2002) provides temporal structure. Every major technological paradigm follows a pattern: installation phase (speculative frenzy, financial capital dominant, growing inequality) $\rightarrow$ turning point (crash, political crisis) $\rightarrow$ deployment phase (productive capital dominant, institutional reform, broadly shared prosperity).

By this model, we are in the installation phase of the AI paradigm. The buildout is speculative, financialized, and concentrating wealth. A turning point is coming—not necessarily a market crash, but a political-economic rupture where the installation-phase institutional arrangements break. What follows the turning point determines whether the technology is deployed productively or destructively.

Perez would say the Marxist-game-theoretic conclusion (it's all multipolar traps, all the way down) is too pessimistic—every prior technological revolution eventually produced a "golden age" of shared prosperity, but only after institutional reform forced by the crisis of the installation phase. The New Deal followed the Gilded Age. Something analogous could follow the AI installation phase. Or it might not—Perez acknowledges the golden age is not guaranteed, just historically precedented.

9.4 Girard: Mimetic Desire and the Scapegoat Mechanism

René Girard's mimetic theory (Girard, 1977, 1986) captures a dimension entirely absent from economic and game-theoretic analysis. Girard argues that desire is not autonomous but imitative: we want what others want, and our sense of our own identity is constructed through mimetic rivalry.

The PMC's consumption was mimetic: the Bay Area house, the consulting career, the credential, the SaaS subscription—these were desired because others desired them. When the PMC is displaced, the crisis is not just economic (less purchasing power) but mimetic (the model of desire that organized professional-class life collapses). The loss of economic position is also a loss of *identity*, and Girard's prediction for what follows identity-crisis at scale is *scapegoating*—the collective designation of a sacrificial victim to channel frustration and restore social cohesion.

The tech oligarchs are the obvious scapegoat candidates. But Girard would note that scapegoating is pre-political, religious in structure, and often misdirected. The current populist relationship to tech moguls is already Girardian: figures like Musk are simultaneously the oligarch (the obvious source of

displacement) and the populist champion (the disruptor of the hated PMC establishment). This is an unstable configuration that will eventually resolve through either elevation (Musk as permanent leader) or sacrifice (Musk as scapegoat when the displacement's effects become undeniable).

9.5 Network Theory and Complexity Economics

The Marxist frame assumes a two-class model (capital vs. labor) that is too simple for the actual dynamics. The real structure is a network: value flows along edges between nodes (firms, workers, institutions, platforms), and the distribution of value is determined by network topology—centrality, clustering, bridge positions, and power-law degree distributions.

The concentration of value in AI is not well-described by "surplus extraction from labor." It is better described by winner-take-all dynamics in scale-free networks: the platform that achieves the most connections captures disproportionate value, regardless of the labor input. This is the preferential attachment mechanism of Barabási and Albert (Barabási and Albert, 1999): nodes that already have many connections attract new connections at a higher rate, producing power-law distributions of connectivity (and therefore of value capture).

The strategic implication is different from the Marxist one. In a Marxist frame, the answer to displacement is "become capital." In a network frame, the answer is "occupy a critical node"—which might mean being a small, well-positioned connector between large systems rather than owning a large system yourself. Integration infrastructure, clearance-mediated access, regulatory expertise—these are *bridge positions* in the network, and bridge positions capture disproportionate value in network economies, not because of capital ownership but because of topological criticality.

10. Synthesis: What Is True

The Marxist frame is the strongest single lens. It identifies the structural compulsion (the flywheel), the mechanism of displacement (rising organic composition of capital), the core contradiction (the realization crisis), and the political economy of the response (control vs. redistribution). Game theory formalizes these dynamics with precision and reveals the nested, self-reinforcing trap structure that makes individual escape rational and collective escape nearly impossible.

The supplementary frames fill specific gaps:

- **Foucault** explains *how* control works at the individual psychological level—not through coercion but through the production of compliant subjects.
- **Polanyi** explains *why* the system might not reach the cybernetic feudalism endpoint—because countermovements are structural, not optional.
- **Perez** provides *temporal structure*—where we are in the cycle and what historically follows.
- **Girard** explains the *cultural and psychological dynamics* of a displaced class—mimetic crisis and scapegoating.

- **Network theory** provides *micro-economic precision* about where value concentrates—not just "capital" but specific topological positions.

No single frame is complete. The Marxist frame is the skeleton. Foucault is the nervous system. Polanyi is the immune system. Perez is the developmental biology. Girard is the psychology. Network theory is the circulatory system. You need the full anatomy to understand the organism.

The one insight that all frames share, stated in Marx's language: **the dynamics are structural, not chosen.** No individual actor—firm, worker, investor, nation, oligarch—is making a mistake. Each is playing optimally given their local information and incentives. The tragedy is that locally optimal play produces globally catastrophic outcomes. The game theory formalizes this. The other frames explain why the globally catastrophic outcome is so difficult to avert even when it is clearly visible in advance.

The window for averting it—the period in which the organizing equilibrium still exists, the counter-movement is still possible, and the surveillance apparatus is not yet comprehensive—is the buildout phase we are currently in. It will not last indefinitely. The flywheel is accelerating under its own power, the traps are interlocking, and the eye is opening.

11. Conclusion: The Moment and the Question

This paper has argued that the historical privilege of the technology worker is ending, and that the mechanism by which it is ending is the same mechanism by which the technology worker was privileged in the first place: the flywheel of capital accumulation that Marx identified in *Capital* (Marx, 1976 [1867]) and that has operated continuously for two and a half centuries. The anomaly being resolved is not labor displacement—that has been continuous since the late eighteenth century—but the exemption of the toolmaker class from the dynamic it served. For roughly fifty years, the small group of people who built the machinery of displacement extracted rents that placed them in the top few percent of the income distribution. Those rents depended on a scarcity—the scarcity of minds capable of translating abstract requirements into working mechanisms—that large language models have now begun to compete away. When the scarcity ends, the rents end, and the toolmaker joins the population whose bargaining power is a function of how easily they can be replaced.

The argument has proceeded through a Marxist frame (Marx, 1976 [1867], 1981 [1894]) augmented by the formal tools of game theory (Nash, 1951; von Neumann and Morgenstern, 1944; Fudenberg and Maskin, 1986). Marx provides the structural compulsion—the fact that individual capitalists do not choose to automate but are forced to by the competitive dynamics of the market. Game theory formalizes that compulsion as a nested system of multipolar traps whose Nash equilibria systematically diverge from their Pareto optima (Alexander, 2014). The central analytical move of the paper is to treat the LLM wave not as a marginal acceleration of the flywheel but as a qualitative change in its dynamics: the tempo of the game is now endogenous to its outcomes, because the machinery of automation has begun to produce the machinery of automation. In the differential equation of Section 4.2, the β term is no longer zero, and the folk-theorem conditions for sustained cooperation (Fudenberg and Maskin, 1986) are being destroyed by the exponential growth of the defector's advantage. The required discount factor δ for cooperative equilibria to exist approaches 1 precisely when firms', nations', and individuals' actual discount factors

are collapsing under competitive urgency.

From this single structural change, the remainder of the argument follows. The realization crisis (Marx, 1981 [1894]; Keynes, 1936)—capital's inability to sell its output to the class whose wages it is destroying—intensifies under the LLM wave because the displaced class is now the high-margin consumer class. The historical patches of credit expansion, government redistribution (Polanyi, 1944), and financialization have diminishing returns when the displaced are no longer weavers or factory workers but the professional-managerial class (Ehrenreich and Ehrenreich, 1977) that sustains premium consumption. The contradiction that Marx identified in the nineteenth century is now operating at the top of the demand curve rather than the bottom, and it is operating at a scale—reasonable estimates place global knowledge-worker compensation between $23 and $70 trillion annually (Forrester, 2018; Manyika et al., 2013; IPPR, 2024)—that dwarfs any prior displacement event.

Capital's response to such displacement has historically been not redistribution but control (Churchill and Vander Wall, 1988; Wagner and Rabuy, 2017; Brennan Center, 2024; McLaughlin et al., 2016), and the current generation of control infrastructure is digital and already being built (Palantir Technologies, 2026; Quartr, 2025; Built In, 2026; American Immigration Council, 2025). The mechanism-design analysis of Section 6 shows why this response is structurally inevitable: a surveillance apparatus that raises the cost of organizing, lowers the probability of its success, and increases the punishment for its failure can eliminate the organizing equilibrium of the displaced class entirely, leaving compliance as the unique Nash outcome. This is not the Pinkertons-and-prisons control of the Gilded Age. It is the Foucauldian control of biopower (Foucault, 1977, 2007): legible, predictive, self-regulating, and cheaper than redistribution by orders of magnitude.

The political realignment of the 2020s is the visible surface of this dynamic. The tech-oligarch–Republican alliance is not an ideological coalition but a structural one, produced by the fact that tech capital and the professional-managerial class are now in a zero-sum game over the addressable displacement market. DOGE (Executive Order 14158, 2025) is the commitment device in the Chicken game between them—an attempt to dismantle the PMC's institutional base within the administrative state before that base can be mobilized to regulate the automation project that the PMC's continued existence would otherwise constrain. The displaced professionals are being defunded, deskilled, and surveilled simultaneously, and the class whose historical function was to administer the state from inside is watching that function be transferred to the tech capital that is displacing it from outside.

Two long-run equilibria remain structurally accessible. The first is what Section 8 called cybernetic feudalism: a stable arrangement in which a small capital-owning class extracts rent from an automated economy, a large surplus population is maintained at subsistence through algorithmic management, and the control apparatus is comprehensive enough to eliminate the organizing threshold. The second is a Polanyian counter-movement (Polanyi, 1944): a cross-class political reaction that re-embeds the economy in social constraints through redistribution, regulation, public ownership, or outright prohibition of certain AI deployments. Polanyi's historical examples—the Factory Acts, the New Deal, the fascist response of the interwar period, the European welfare state—show that counter-movements are structurally near-inevitable when disembedding reaches a critical threshold, but they do not show which form the counter-movement will take, and they include outcomes ranging from democratic redistribution to authoritarian consolidation.

Perez (Perez, 2002) provides additional temporal structure: the installation phase of every technological revolution ends in rupture, and the deployment phase that follows is sometimes a golden age and sometimes a catastrophe. The outcome is not determined by the technology itself.

Which equilibrium obtains depends on a race condition. The rate at which the control apparatus is deployed must be compared to the rate at which political mobilization of the displaced class reaches its critical threshold. The buildout phase—roughly 2025 through 2032—is the window in which that race is being run. After it, the path-dependent dynamics of the surveillance-adjusted coordination game are likely to lock in one outcome or the other, and the options available to the displaced will be the options the control architecture of that moment permits.

The paper does not predict which equilibrium will obtain, because no single frame is adequate to the prediction and the synthesis of frames presented in Section 10 produces a system of dynamics too complex to admit confident forecasting. The Marxist frame captures the structural compulsion but underweights political contingency; the Foucauldian frame captures the micro-mechanisms of control (Foucault, 1977, 2007) but treats power as more diffuse than the game-theoretic analysis requires; the Polanyian frame captures the counter-movement's inevitability without specifying its form (Polanyi, 1944); the Perezian frame captures the technological-cycle timing (Perez, 2002) but is silent on the specific institutional arrangements of the deployment phase; the Girardian frame captures the scapegoating dynamics of displaced classes (Girard, 1977, 1986) but is pre-political and often misdirected; the network-theoretic frame captures the topological concentration of value (Barabási and Albert, 1999) but treats class as an emergent rather than structural category. What all frames agree on is that the dynamics are structural rather than chosen, that locally optimal play produces globally catastrophic outcomes, and that the window for averting the catastrophic outcome is open but closing.

The final claim of the paper is not a prediction but a naming. What is happening to the technology worker is formally the same thing that happened to the handloom weaver, with one critical difference: the weaver did not also build the loom. For two and a half centuries, the flywheel required a privileged cognitive class to operate, and that class was compensated accordingly. The privilege is ending now, not because technology workers have done anything wrong, but because the thing they built has finally become competent enough to reproduce itself. The class that spent fifty years building the flywheel is becoming the class on which the flywheel acts. This is not a tragedy in any moral sense—the flywheel is amoral, and capital is structurally compelled toward it—but it is a tragedy in the older, theatrical sense: a specific character arrives at a specific end through the operation of forces the character itself set in motion.

The open question is whether that character, understanding the forces in play, can coordinate with the other characters in the drama before the architecture of the stage itself is rebuilt to prevent coordination. The paper has tried to describe the forces with enough precision that the question can at least be asked clearly. The answer is a matter for the next decade, not the next paper, and it will be settled by political action or the absence of it rather than by further analysis. The window is open. It is closing. The eye is opening. What happens next is not yet written.

References

Alexander, Scott [Scott Siskind]. 2014. "Meditations on Moloch." *Slate Star Codex*, July 30, 2014. https://slatestarcodex.com/2014/07/30/meditations-on-moloch/.

American Immigration Council. 2025. "ICE to Use ImmigrationOS by Palantir, a New AI System, to Track Immigrants' Movements." Published April 2025. https://www.americanimmigrationcouncil.org/blog/ice-immigrationos-palantir-ai-track-immigrants/.

Barabási, Albert-László and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286 (5439): 509–512.

Brennan Center for Justice. 2024. "Social and Economic Harm." End Mass Incarceration initiative. Accessed 2026. https://www.brennancenter.org/issues/end-mass-incarceration/social-and-economic-harm.

Rowe, Lisa. 2026. "What Is Palantir? The Company Behind Government AI Tools." *Built In*, February 2026. https://builtin.com/articles/what-is-palantir.

Churchill, Ward and Jim Vander Wall. 1988. *The COINTELPRO Papers: Documents from the FBI's Secret Wars Against Dissent in the United States*. Boston: South End Press.

Ehrenreich, Barbara and John Ehrenreich. 1977. "The Professional-Managerial Class." *Radical America* 11 (2): 7–31.

Executive Order 14158. 2025. "Establishing and Implementing the President's Department of Government Efficiency." Federal Register 90: 8441. January 20, 2025.

Gillett, Frank. 2018. "The Global Information Worker Population Swells to 1.25 Billion in 2018." Forrester Research. https://www.forrester.com/blogs/the-global-information-worker-population-swells-to-1-25-billion-in-2018/.

Foucault, Michel. 1977. *Discipline and Punish: The Birth of the Prison*. Translated by Alan Sheridan. New York: Vintage Books. (Originally published 1975.)

Foucault, Michel. 2007. *Security, Territory, Population: Lectures at the Collège de France, 1977–1978*. Edited by Michel Senellart, translated by Graham Burchell. New York: Palgrave Macmillan.

Fudenberg, Drew and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54 (3): 533–554.

Girard, René. 1977. *Violence and the Sacred*. Translated by Patrick Gregory. Baltimore: Johns Hopkins University Press. (Originally published 1972.)

Girard, René. 1986. *The Scapegoat*. Translated by Yvonne Freccero. Baltimore: Johns Hopkins University Press.

Institute for Public Policy Research. 2024. *Transformed by AI: How Generative Artificial Intelligence Could Affect Work in the UK*. London: IPPR.

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.

Manyika, James, Michael Chui, Jacques Bughin, Richard Dobbs, Peter Bisson, and Alex Marrs. 2013. *Disruptive Technologies: Advances That Will Transform Life, Business, and the Global Economy*. San Francisco: McKinsey Global Institute.

Marx, Karl. 1976 [1867]. *Capital: A Critique of Political Economy, Volume I*. Translated by Ben Fowkes. London: Penguin Classics.

Marx, Karl. 1981 [1894]. *Capital: A Critique of Political Economy, Volume III*. Translated by David Fernbach. London: Penguin Classics.

McLaughlin, Michael, Carrie Pettus-Davis, Derek Brown, Chris Veeh, and Tanya Renn. 2016. *The Economic Burden of Incarceration in the U.S.* Working Paper CI072016. Institute for Advancing Justice Research and Innovation, Washington University in St. Louis.

Nash, John F. 1951. "Non-Cooperative Games." *Annals of Mathematics* 54 (2): 286–295.

"Palantir Technologies." 2026. *Wikipedia*. Accessed February 2026. `https://en.wikipedia.org/wiki/Palantir_Technologies`.

Perez, Carlota. 2002. *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Cheltenham, UK: Edward Elgar.

Polanyi, Karl. 1944. *The Great Transformation: The Political and Economic Origins of Our Time*. Boston: Beacon Press.

Quartr. 2025. "Palantir: Turning Data into Decisions." September 2025. `https://quartr.com/insights/company-research/palantir-turning-data-into-decisions`.

TECHi. 2026. "Anthropic IPO (Claude AI): $60B Raise, Valuation & How to Invest 2026." April 2026. `https://www.techi.com/anthropic-ipo/`.

von Neumann, John and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Wagner, Peter and Bernadette Rabuy. 2017. "Following the Money of Mass Incarceration." Prison Policy Initiative, January 2017. `https://www.prisonpolicy.org/reports/money.html`.